

Topology Design for Directional Range Extension Networks with Antenna Blockage

Thomas Shake
MIT Lincoln Laboratory
shake@ll.mit.edu

Abstract—Extending the range of local area surface networks by using small aircraft to relay traffic to geographically distant areas is a frequently considered topic in military network technology development. This paper considers the use of a modular pod design for incorporating highly directional antennas and associated electronics into small aircraft to perform such range extension. In particular, the paper examines trade-offs in network topology design introduced by pod-based antenna blockages. Using certain modeling approximations, the paper presents a quantitative analysis showing design trade-offs among several design parameters, including the number of antenna beams supported by the pod design, the number of surface nodes to be supported by each aircraft, and the topology characteristics of the aerial relay network. The analysis suggests that low-degree air topologies such as rings and strings be used to maximize the connection availability of surface nodes, and to optimize a trade-off between connection availability and the total number of surface nodes that can be connected to the backbone (and hence to each other). The analysis also results in an optimization criterion for the formation and maintenance of low-degree air topologies.¹

I. INTRODUCTION

Tactical military networks both on land and at sea often have restricted transmission ranges due to limits on terminal transmission power, geographic features that block line-of-sight, and poor over-the-horizon signal propagation. Limiting communications to local tactical areas can strongly constrain mission success, and some form of range extension is often needed for these tactical networks. One frequently considered option for range extension is the use of relatively small unmanned aircraft to serve as over-the-horizon relay platforms connecting geographically separated local groups of surface nodes. Fig. 1 illustrates a conceptual example of such a range-extension scenario that includes both maritime and land-based network terminals.




Fig. 1. Conceptual example of network range extension via aerial platforms

¹This work is sponsored by the United States Department of the Navy under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, recommendations and conclusions are those of the authors and are not necessarily endorsed by the United States Government.

A promising method of incorporating electronics and antennas into an airborne relay platform is to enclose them in a separate pod, which can be modularly designed and attached to the bottom of the wings or fuselage of a variety of aircraft types [1]. Fig. 2 shows a generic concept for such a pod design with potential antenna placements on each end of the pod and on the bottom. There is great advantage to making these antennas highly directional; this enables both a much higher data rate connection between aircraft and a higher level of rejection of local interference in the surface-air connections. Such directional antennas may be implemented by mechanically steered dish antennas or by electronically steered devices (e.g., phased arrays). However, using highly directional antennas also makes the design, discovery, and maintenance of the relay network topology significantly more difficult.




Fig. 2. Generic pod concept for modular aerial electronics

Depending on which of the potential antenna placements shown in Fig. 2 is used for a given design, the line-of-sight (LOS) between two aircraft or between aircraft and surface nodes may undergo intermittent blockages as aircraft fly in patterns that enable them to maintain LOS to a local surface node group. (The flightpaths typically considered for such range extension purposes are closed paths with circular or racetrack-shaped patterns [1].) Antennas placed on the ends of the pod are especially vulnerable to regular blockages, and such blockages are an important factor affecting the design and management of the network topology [2][3]. In particular, such blockages can cause significant periods of unavailability of the range extension capability as seen by surface nodes in local groups. This paper considers the situation where the relay network antennas are placed on either end of the pod, but not on the bottom. This may correspond to design scenarios where the pod physical design does not allow placement of directional antennas on the bottom, or to scenarios where antennas on the pod bottom must be reserved for other uses

such as sensor suites or other communication systems.

This paper derives mathematical relationships among various design parameters and performance measures using certain modeling approximations appropriate to aerial relay networks. These relationships are used to quantify trade-offs among relay network availability, the degree of inter-connectivity that can be attained among surface nodes in separated local groups, the number of surface nodes per aircraft that can be accommodated, and the number of directional antenna beams that can be incorporated into a given pod design.

The outline of the rest of the paper is as follows. Section II surveys previous work in this area. Section III describes the range extension network model and design parameters. Section IV develops quantitative relationships among key design elements and performance metrics. Section V considers some implications of the results in Section IV for range extension network topology design and for technology development. Section VI summarizes the results and offers some concluding thoughts.

II. PREVIOUS WORK

Topology management for wireless ad hoc and sensor networks is an area that has received extensive attention in the research literature. An extensive survey is contained in [4]. However, the great majority of this work assumes either static nodes, omni-directional transmissions, or both. There is also a significant body of work on Medium Access Control (MAC) for wireless ad hoc networks with directional antennas, which is surveyed in [5], and which often contains material on topology discovery and/or control. Other work in modeling topology management with directional aircraft antennas has been presented in [6] and [7] for radio frequency transmission, and [8]-[10] for optical free-space networks. None of this work, as far as the author is aware, deals with antenna blockages due to structural factors such as aircraft fuselage or pod structures. The concept of pod-based antenna blockage has been brought up in [1]-[3]. The first of these papers examines the effects of platform dynamics and aircraft orbit shapes on connection availability, but does not generalize the analysis to arbitrary topologies or consider more than two antennas per pod. The second and third of these papers concentrate on the effects of antenna switching on specific small-scale scenarios. The current work models pod-based antenna blockages in a way that allows generalization to arbitrary topologies with arbitrary numbers of aircraft and surface nodes, and presents new closed-form analytical results for low-degree air topologies.

III. RANGE EXTENSION TOPOLOGY MODEL

A. Modeling Parameters

This paper explores trade-offs and optimizations in topology design for the types of range extension networks described qualitatively above. To quantify these trades and optimizations, the following formal model is established. Figs. 3 and 4 illustrate the model of the pod design and range extension network topology considered here. Each relay aircraft is assumed to

have one pod which has n narrow antenna beams, $n/2$ of which are placed on each end of the pod. The relay network topology model consists of N aircraft, each of which supports S surface nodes. Each surface node has exactly one antenna beam which can be pointed at the nearest aircraft for range extension. Aircraft nodes do not source or sink traffic, but simply relay network traffic among surface nodes. (Surface nodes under a single aircraft may or may not have other means of local communication that do not involve the aircraft relay.) The aircraft are assumed to be relatively far away from each other (otherwise the range extension provided would be poor), and to fly in approximately circular flightpaths with radii that are small compared to the distance between the aircraft. A representative node lay-down with three example connection topologies is illustrated in Fig. 5. This figure shows the centroids of racetrack flightpaths as light circles; the realistically modeled racetrack flightpaths are the green lines around the centroids. The surface nodes are shown with linear paths, and the light triangles represent specific positions of the nodes at a particular point in time.




Fig. 3. Pod antenna model




Fig. 4. Abstract topology model for range extension network

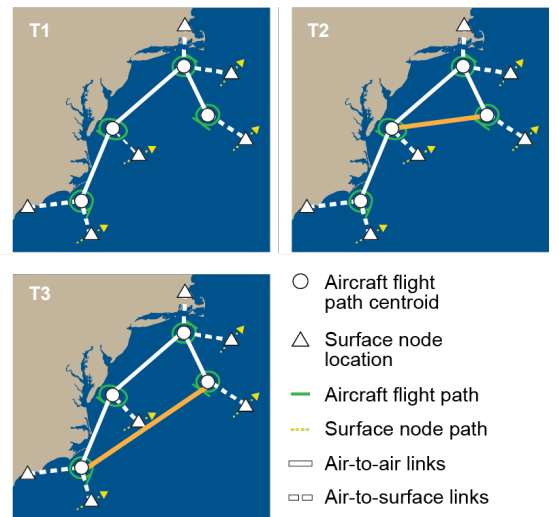


Fig. 5. Representative practical range extension node lay-down with alternative topologies

B. Key Modeling Approximations

The antenna blockage modeling in this paper is based on two key approximations:

- 1) Compass directions from a given air node to other nodes are assumed to be approximately constant as the aircraft moves.
- 2) Antenna Field-of-View (FOV) blockages are modeled by fixed azimuth and elevation limits.

These approximations enable a general mathematical analysis of an otherwise specific and complicated situation. The first approximation is that the relative compass direction of both air and surface nodes as seen from any air node changes only negligibly during the time it takes to complete one circuit of the closed flightpath. This should be approximately true for other air nodes for most practical range extension scenarios. It should also be approximately true for many surface node distributions, but will not be such a good approximation for scenarios where the surface nodes are located close to the air relay node.

The second key approximation is in how antenna field of view (FOV) on each pod is modeled. Since all air nodes are likely to fly at approximately the same altitude, the required elevation angle (in aircraft centric coordinates) to other aircraft will stay within a narrow range which is mostly free of blockages caused by elevation angle restrictions. Due to likely pod geometries, shown generically in Figs. 2 and 3, the main cause of blockage in the azimuth component will be due to the pod body. This paper approximates the FOV of the antennas on each end of the pod as being clear for an azimuthal angle of $(180+2\gamma)^\circ$ around antenna boresight, and approximates the elevation FOV as being limited to δ° past nadir, as illustrated in Fig. 6. This approximation should be fairly accurate for the type of pod design considered and for FOV in the directions of most air and surface node locations. Aircraft wing blockages in particular are neglected in this approximation, but wing blockages generally represent relatively short periods of blockage for circular and racetrack flightpaths, so this approximation should not have a great deal of effect on the line-of-sight availability of links averaged over the time it takes an aircraft to complete one circuit of its flightpath.

This paper presents quantitative results for the case of $\gamma = 0$ and $\delta = 0$ (representing no overlap between the FOV of the antennas on opposite ends of the pod). A subsequent paper will present results in the more complex case of $\gamma > 0$ and $\delta > 0$.

Given the approximations just described, the connections of a specific air node to other nodes (both air and surface) can be modeled by the azimuth angle of the pointing vector to each of the nodes that are visible and within transmission range of the aircraft in question. This is illustrated in Fig. 7, which shows the pointing vectors to hypothetical air and surface nodes in the aircraft-centric azimuth of air node i . Since the flightpath is approximated by a circle of essentially zero radius, as the aircraft flies around one circuit of its path (which takes time




Fig. 6. Approximate antenna field-of-view model

T^i for the i^{th} air node) the directions to the other nodes simply rotate in the azimuth plane of air node i . Fig. 7 also defines angular directions β to the other air nodes j, k, l . One final parameter of interest is the number of connections of air node i to other air nodes (i.e., the degree of air node i with respect to the air-to-air topology of the relay network), which will be denoted by d_a^i .




Fig. 7. Azimuth plane view of air node i

IV. TOPOLOGY DESIGN TRADE-OFFS

A. Design Criteria

This paper focuses on maximizing both the total number of surface nodes that can be connected to the air backbone (and hence to each other) and the availability of these connections.² The performance metric for the number of connectable surface nodes will be denoted by $\langle K \rangle / N$, the average number of surface nodes per aircraft that can be supported by the air topology. The figure of merit used here for the availability of these connections will be the average fractional connection availability of the entire ensemble of surface nodes in the network, which is given by $A_f = \langle K \rangle / (N \times S)$. (This may be apportioned either evenly or unevenly among surface nodes.)

The metric $\langle K \rangle / N$ depends on several factors in this network model, including the number of antenna beams (presumably restricted by hardware design constraints and possibly by processing constraints for phased-array antennas); the number and compass direction of air backbone nodes (determined by the locations of the aircraft and by the air topology design);

²Other optimizations were considered, such as maximizing some form of all-to-all traffic, or optimizing the survivability of the air backbone. The first is a complex problem involving multi-commodity flow optimization [11] and fairness issues [12], and the second seems poorly suited to this type of network since the restriction that each surface node to a single antenna sets the maximum connectivity of surface-to-air nodes at 1.

and the number and locations of the surface nodes relative both to their local aircraft and to the compass directions of the local aircraft's connections to other aircraft.

B. General results

Given the problem formulation above with $\gamma = \delta = 0^\circ$, there are a few simple and general results which are readily apparent. There is an upper bound to the number of surface nodes per air node that can be connected, which is given by

$$\frac{K^U}{N} = \frac{1}{N} \sum_{i=1}^N (n - d_a^i) \quad (1)$$

This is a straightforward consequence of the fact that antenna beams that are used for the air topology are not available to be used to connect surface nodes. The average number of connected surface nodes per air node, $\langle K \rangle / N$, can approach this upper bound either by the optimum choice of specific fixed locations for all surface nodes or by simply increasing S , the number of surface nodes underneath each aircraft. Increasing S , however, may lower the overall average fractional node connection availability and may prevent the upper bound in (1) from being attainable. If one wishes to guarantee that all surface nodes have 100% connection availability ($\langle A_f \rangle = 1$), the following criteria will do that under the appropriate conditions. If surface node locations are arbitrary and unknown, 100% average connection availability can be attained only when

$$S \leq \frac{n}{2} - \max_i(d_a^i) \quad (2)$$

which guarantees that the upper bound of (1) is not attained. If surface node locations are optimally placed in fixed locations, $\langle A_f \rangle = 1$ can be attained only when

$$S \leq n - \max_i(d_a^i) \quad (3)$$

with the upper bound K^U/N being attained when equality holds in (3).

Fig. 8 shows the relationship between the number of surface nodes, S , that can be supported per aircraft and the number of antenna beams per pod, n , while guaranteeing $\langle A_f \rangle = 1$. The solid lines in the plot represent the relationship for randomly located surface nodes, and the dotted lines represent the relationship for optimally located, fixed surface node locations. In addition, the relationship is plotted for two different values of the maximum degree of the air nodes in the overall air topology $\max_i(d_a^i)$, the blue curves being for $\max_i(d_a^i) = 2$ and the red curves for $\max_i(d_a^i) = 3$.

Fig. 9 shows the maximum degree of the air nodes that can be attained while still maintaining $\langle A_f \rangle = 1$ for $S = 1$ surface node per aircraft as a function of the number of antenna beams per pod, n . This plot indicates a trade-off between surface node availability and capacity and/or survivability for the case when a single surface node is to be supported under each aircraft. Both the capacity and survivability of the air backbone are related to the maximum degree of the air nodes, though not




Fig. 8. Minimum number of antenna beams, n , required for $\langle A_f \rangle = 1$

in straightforward ways. Low air node degrees certainly limit both air backbone capacity and survivability, but higher node degrees do not necessarily guarantee either higher capacity or better survivability. Still, this figure does illustrate a general relationship between maintaining full surface node availability and a minimum capability to provide air backbone capacity and survivability.




Fig. 9. Maximum air node degree for $S = 1$ and $\langle A_f \rangle = 1$

C. Ring and String Air Topologies

The results in Section IV-B imply that the degree of the nodes in the air topology should be limited to maximize the interconnectivity of surface nodes and their connection availability. This suggests that ring or string topologies should strongly be considered for the air backbone. In a ring topology, $d_a^i = 2$ for all air nodes; for a string, $d_a^i = 2$ for all nodes except the ones on the ends of the string, where $d_a^i = 1$. With these restrictions, and for the case of $\gamma = 0$ and $\delta = 0$, the value of $\langle K \rangle / N$ can be written as a closed form equation. This equation can be derived by considering each single air node in isolation. For a given node i , define time t_0 such that one of the neighbor air nodes (the only one in the case of a string end node) is positioned directly at the nose of the pod, and such that the other neighbor node (if any) is at an azimuth angle β_2 such that $\beta_2 < 180^\circ$, as illustrated in Fig. 10. At time t_0 , surface nodes may be located at azimuth angles that are either known and specified, or are unknown and assumed random. The case of randomly located surface nodes is considered first.




Fig. 10. Node locations from air node i for ring or string topologies

1) *Randomly Located Surface Nodes*: This analysis assumes that surface nodes of unknown location are equally likely to be on either the "North" (pod nose) or "South" (pod tail) sides of the pod. In this case, at t_0 the surface nodes are binomially distributed, with binomial parameters $p = 0.5$ and $N_b = S$. The number of surface nodes that can be supported at time t_0 can then easily be calculated. This number can then be averaged over one complete circuit of the aircraft's flightpath (taking time T^i for the i^{th} air node). Given the approximations of Section III, one complete circuit of node i 's flightpath is equivalent to a 360° rotation of the pod's azimuth orientation. The average number of connected surface nodes per air node in the entire topology is thus given by

$$\begin{aligned} \frac{\langle K \rangle}{N} = \frac{1}{N} \sum_{i=1}^N \left\{ \frac{\beta_2^i}{180} \left[\sum_{j=0}^S \frac{1}{2^S} \binom{S}{j} \left(\min\left[\frac{n}{2}, (d_a^i - 1 + j)\right] \right. \right. \right. \\ \left. \left. \left. + \min\left[\frac{n}{2}, (S - j + 1)\right] \right) \right] + \left(1 - \frac{\beta_2^i}{180} \right) \left[\sum_{j=0}^S \frac{1}{2^S} \binom{S}{j} \right. \right. \\ \left. \left. \times \left(\min\left[\frac{n}{2}, j\right] + \min\left[\frac{n}{2}, (S - j + d_a^i)\right] \right) \right] - d_a^i \right\} \quad (4) \end{aligned}$$

and the total average connection availability of surface nodes in the network is

$$\langle A_f \rangle = \frac{\langle K \rangle}{N \times S} \quad (5)$$

Fig. 11 shows some sample results of these calculations for two simple, regular air topologies. The examples use $n = 4$ antenna beams per aircraft pod, with two on each end of the pod. The top plots in each case show the total average number of connected surface nodes per air node as a function of the number of surface nodes S under each aircraft, with the blue curve showing results for the string topology and the red curve showing results for the ring topology. The upper bounds for each topology are shown by the red and blue dotted lines. The bottom plots in each case show the average connection availability for surface nodes. The string topologies outperform the ring topologies for both metrics for any given value of S , which is expected since the string topologies have more total antennas available to be used for surface connections. Note that there is a trade-off between maximizing the total average number of connected surface nodes and the average availability of each surface node's connection to the air backbone. Also, in the case shown in these particular examples, the average connection availability $\langle A_f \rangle$ is never 100% because there are too few antennas to meet the criteria of (2).

2) *Known Surface Node Locations*: A similar procedure can be used to calculate the average number of connected

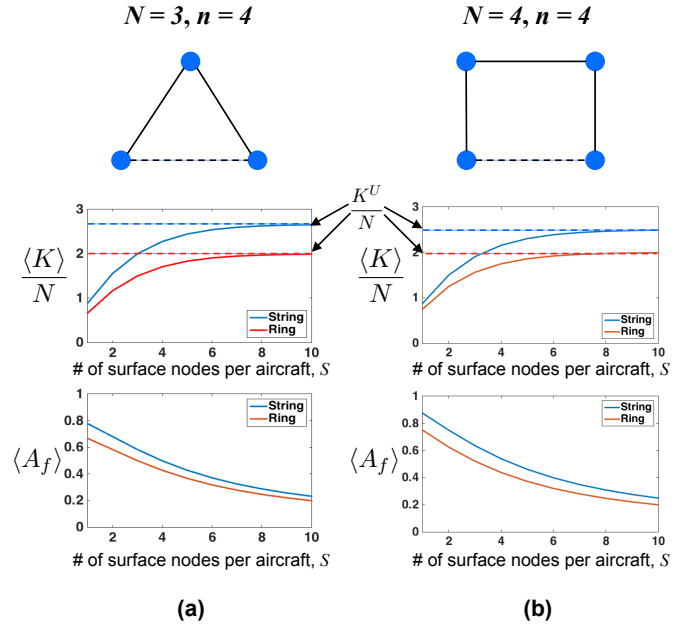


Fig. 11. Expected connectivity and availability for two simple examples

surface nodes and their average availability when the locations of the surface nodes are known, assuming that these locations are either fixed or approximately fixed for the duration, T^i , of one flightpath circuit. Let f_j be defined as the fraction of time T^i that j surface nodes can be connected. The expected number of surface nodes connected to node i is given by

$$\langle K_F^i \rangle = \sum_{j=0}^S j \times f_j \quad (6)$$

and the expected total number of connected surface nodes per air node is

$$\frac{\langle K_F \rangle}{N} = \frac{1}{N} \sum_{i=1}^N \langle K_F^i \rangle \quad (7)$$

The calculations for f_j are not simple to write in closed form for most cases of interest, but can be easily calculated by a simple heuristic.³ The heuristic can also incorporate a mixture of known and unknown (binomially distributed) surface node locations in the computations of expected connectivity and availability, and accommodates different numbers of surface nodes under each air node (i.e., $S_i \neq S_j$).

Some sample results for the three range extension topologies of Fig. 5 are shown in Table I. The top left topology (T1) in Fig. 5 is a string topology, while the top right and lower left topologies (T2 and T3 respectively) augment this topology by adding one air-air link in each case. Topology T3 represents a ring (though it may be impractical in the example illustrated due to the long range between the Southernmost aircraft and the Easternmost aircraft). The results in Table I show the expected average connectivity and availability for each topology

³The heuristic is too long to be included in this paper but will be described in a subsequent paper mentioned above including results for $0 < \gamma < 90^\circ$.

for two different values of n and for two separate cases: a) fixed locations of all land and sea surface nodes; and b) fixed land node locations and random sea node locations. Note that for $n = 4$ antenna beams, the air topology augmentations decrease both the expected surface node connectivity and the expected surface node availability. If n is increased to six, however, the decrease in both metrics is minimal, and likely worth the trade-off to get extra connectivity in the air backbone.

	$n = 4$ antenna beams			$n = 6$ antenna beams		
	T1 (string)	T2	T3 (ring)	T1 (string)	T2	T3 (ring)
Fixed land and ship node locations						
$\langle K \rangle / N$	1.35	1.19	1.29	1.5	1.46	1.5
$\langle A_f \rangle$	0.9	0.79	0.86	1	0.98	1
Fixed land, random ship node locations						
$\langle K \rangle / N$	1.31	1.10	1.14	1.5	1.48	1.48
$\langle A_f \rangle$	0.88	0.73	0.76	1	0.99	0.99

TABLE I

EXPECTED CONNECTIVITY AND AVAILABILITY FOR EXAMPLES IN FIG. 5

V. DESIGN IMPLICATIONS

A. Topology Algorithm Selection

The sample results in previous sections indicate that forming and maintaining ring or string air topologies yields the best combination of surface node connection availability and number of surface node connections that can be supported. Depending on the design constraints on the number of antenna beams that a pod can support, a string topology formation and maintenance algorithm may be the best design choice for small scale aerial range extension networks. As the number of air nodes increases, the differences in connectivity and availability between ring and string topologies decrease (compare Fig. 12 with Fig. 11), but many practical range extension networks may contain just a few air nodes. A careful examination of (4) reveals that the availability of surface node connections is maximized by maximizing the sum $\sum_{i=1}^N \beta_2^i$. This is due to the fact that for each air node i , the sum in (4) that is multiplied by $(\beta_2^i/180)$ is always greater than or equal to the sum that is multiplied by $(1 - \beta_2^i/180)$. Fig. 13 illustrates this criterion for topology optimization for a set of six arbitrarily placed nodes. This maximization is a computationally difficult problem, however, though it may be practical to solve this for very small scale topologies. A more computationally efficient approach would be to form and maintain a Minimum Weight Spanning Tree (MST) topology [13] for the air nodes. However, for $N > 3$ a minimum weight spanning tree topology does not always maximize surface node availability, and will sometimes include nodes of degree > 2 . A degree constrained MST could be calculated, but this is equivalent to the Hamiltonian Path Problem, which is NP-hard. More work on computationally efficient topology formation and maintenance algorithms for this scenario is needed.

B. Technology Development Considerations

In the near term, small values of n may be practical with mechanically steered or phased array antenna designs. Size,




Fig. 12. Expected connectivity and availability with a sample larger air topology




Fig. 13. Illustration of the topology optimization criterion of Section V-A

weight, and power requirements are key considerations for a pod design that must be mounted to a small aircraft, and these factors are likely to strongly constrain n in the near term. Small values of n do place stronger constraints on the connectivity of the air backbone (to keep surface node availability high) and also limit the survivability of the air backbone. The ability to increase practical range extension network availability by performing make-before-break antenna switches on air-air links is also constrained by having small n . Larger values of n allow the support of greater numbers of surface nodes for a given level of air backbone connectivity, and may also allow simpler topology management techniques to be used for forming and maintaining the air topology. (One reason for this is that MST algorithms are likely to result in a smaller penalty in surface node connection availability with large n .) Therefore, developing technologies that can increase the potential value of n with minimum size, weight, and power requirements should be a priority for increasing the performance of aerial range extension networks.

Increasing the number of antennas available on the pod may result in greater complexity in other ways that have not been the focus of this paper. For example, frequency and spectrum management will almost certainly be more complicated with large n since more beams must be managed, and receiver front end implementations may have significant levels of signal crosstalk between processing chains for each beam, making frequency reuse more difficult. Scheduling transmissions and receptions in networks with large numbers of directional beams is also a potential issue. In the far term, digital beamforming approaches with phased array antennas may offer a solution to this potential problem [14], but in the

near term it may be problematic.

C. Future Work

The work in this paper can be both generalized to reduce the specificity of some of the assumptions and approximations, for example extending the calculations to less restrictive field-of-view assumptions. In particular, further work is already in progress to extend the calculations to overlapping antenna fields-of-view from each end of the pod (i.e., $0 < \gamma \leq 90^\circ$). The calculations can also be made more accurate for specific effects, including incorporation of wing blockages, non-zero switching times when antennas must be switched from one destination node to another, and effects of network layer effects such as routing reachability on the availability of end-to-end connections.

VI. CONCLUSIONS

This paper has explored topology-related design trade-offs for aerial range extension networks using highly directional antennas on relay aircraft. The paper presented a design in which aircraft electronics and antenna structures are modularly encapsulated in a pod attaching to the underside of the aircraft, and incorporating multiple antenna beams placed on each end of the pod. The paper developed approximate models of the effects of antenna beam blockages resulting from representative pod geometries in typical range extension network topologies. The primary figures of merit considered in this paper are the average number of surface nodes that can be connected to the aerial range-extension backbone, and the average availability of the connections from these surface nodes to the backbone, and hence to each other. Both of these factors are constrained by the antenna blockage characteristics of the pod-based design; the paper quantified how these factors may be traded against one another as a function of the total number of aircraft, the aircraft topology, the number of surface nodes supported per aircraft, and the number of antenna beams supplied by the pod design.

The number of antenna beams that can be designed into the aircraft electronics pod is a primary factor limiting both the number of surface nodes that can be connected and their connection availability. For example, increasing this number from four to six beams per aircraft allowed the average connection availability of six total surface nodes in a representative four-aircraft air topology to be increased from 76% to 100% while maintaining a 2-connected air backbone.

For small numbers of antenna beams per pod (e.g., 4-6) the results of the analysis suggest that low-degree air topologies (either a ring or string) be the goal of air topology formation and maintenance algorithms to maximize the connection availability of surface nodes to the air backbone. The paper has derived a criterion for air topology formation based on the compass directions of neighboring air nodes, and recommends further research into computationally efficient algorithms that optimize this criterion.

ACKNOWLEDGEMENTS

The author would like to thank Terrance Gibbons of MIT Lincoln Laboratory's Tactical Networks Group for supplying details of the example four-aircraft scenarios in Fig. 5.

REFERENCES

- [1] D. Mehta, B. Ganguly, "The Effect of Platform Dynamics on Aerial Layer Network Performance", IEEE Milcom2014, Oct. 2014
- [2] J. Wang, P. Deutsch, A. Coyle, T. Shake, B-N Cheng, "An Implementation of a Flexible Topology Management System for Aerial High Capacity Directional Networks", IEEE Milcom2015, Oct. 2015
- [3] J. Wang, T. Shake, P. Deutsch, A. Coyle, B-N Cheng, "Topology Management Algorithms for Large-Scale Aerial High Capacity Directional Networks", submitted to IEEE Milcom2016, Nov. 2016
- [4] P. Santi, "Topology Control in Wireless Ad Hoc and Sensor Networks", *ACM Computing Surveys*, Vol. 37, No. 2, June 2005
- [5] O. Bazan, M. Jaseemuddin, "A Survey On MAC Protocols for Wireless Adhoc Networks with Beamforming Antennas", *textslIEEE Communication Surveys and Tutorials*, Vol. 14, No. 2, Second Quarter 2012
- [6] D. Van Hook, M. Yeager, J. Laird, "Automated Topology Control for Wideband Directional Links in Airborne Military Networks", IEEE Milcom2005, Oct. 2005
- [7] P. Wang, B. Henz, "Antenna Assignment for JALN HCB", IEEE Milcom2015, Oct 2015
- [8] J. Zhuang, M. Casey, S. Milner, S. Gabriel, G. Baecher, "Multi-Objective Optimization Techniques in Topology Control of Free Space Optical Networks" IEEE Milcom2004.
- [9] A. Desai, S. Milner, "Autonomous Reconfiguration in free-Space Optical Sensor Networks", *textslIEEE Journal on Selected Areas in Communications*, Vol. 23, No. 8, Aug. 2005
- [10] A. Kashyap, K. Lee, M. Kalantari, S. Khuller, M. Shayman, "Integrated topology control and routing in wireless optical mesh networks", *Computer Networks*, Vol 51, pp. 4237-4251, 2007
- [11] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms (3rd ed.)*. MIT Press, 2009
- [12] T. Bonald, L. Massoulié, "Impact of fairness of Internet performance", ACM SIGMETRICS 2001
- [13] D. Bertsekas, R. Gallager, *Data Networks*, 2nd ed., Prentice Hall, 1987
- [14] G. Kuperman, R. Margolies, N. Jones, B. Proulx, A. Narula-Tam, "Uncoordinated MAC for Adaptive Multi-Beam Directional Networks: Analysis and Evaluation", submitted to ICCCN 2016